

APPLICATION
FOR
UNITED STATES LETTERS PATENT

INTERNATIONAL BUSINESS MACHINES CORPORATION

GB920000100US1 - CPA
Borman et al

A COMPUTER SYSTEM WITH TWO HEAPS IN CONTIGUOUS STORAGE

BACKGROUND OF THE INVENTION

Field of the Invention

The invention relates to a computer system supporting an object-oriented environment having storage, at least a portion of which is divided into multiple heaps.

Description of the Related Art

Programs written in the Java programming language (Java is a trademark of Sun Microsystems Inc) are generally run in a virtual machine environment, rather than directly on hardware. Thus a Java program is typically compiled into byte-code form, and then interpreted by a Java virtual machine (JVM) into hardware commands for the platform on which the JVM is executing. The JVM itself is an application running on the underlying operating system. An important advantage of this approach is that Java applications can run on a very wide range of platforms, providing of course that a JVM is available for each platform.

Java is an object-oriented language. Thus a Java program is formed from a set of class files having methods that represent sequences of instructions (somewhat akin to subroutines). A hierarchy of classes can be defined, with each class inheriting properties (including methods) from those classes which are above it in the hierarchy. For any given class in the hierarchy, its descendants (i.e. below it) are call subclasses, whilst its ancestors (i.e. above it) are called

superclasses. At run-time objects are created as instantiations of these class files, and indeed the class files themselves are effectively loaded as objects. One Java object can call a method in another Java object. In recent years Java has become very popular, and is described in many books, for example "Exploring Java" by Niemeyer and Peck, O'Reilly & Associates, 1996, USA, and "The Java Virtual Machine Specification" by Lindholm and Yellin, Addison-Wedley, 1997, USA.

The standard JVM architecture is generally designed to run only a single application, although this can be multi-threaded. In a server environment used for database transactions and such-like, each transaction is typically performed as a separate application, rather than as different threads within an application. This is to ensure that every transaction starts with the JVM in a clean state. In other words, a new JVM is started for each transaction (i.e. for each new Java application). Unfortunately however this results in an initial delay in running the application (the reasons for this will be described in more detail later). The overhead due to this frequent starting and then stopping a JVM as successive transactions are processed is significant, and seriously degrades the scalability of Java server solutions.

Various attempts have been made to mitigate this problem. EP-962860-A describes a process whereby one JVM can fork into a parent and a child process, this being quicker than setting up a fresh JVM. The ability to run multiple processes in a Java-like system, thereby reducing overhead per application, is described in "Processes in KaffeOS: Isolation, Resource Management, and Sharing in Java" by G back, W Hsieh, and J Lepreau

(see:

<http://www.cs.utah.edu/flux/papers/kaffeos-osdi00/main.html>).

5 Another approach is described in "Oracle JServer Scalability and Performance" by Jeremy Litz, July 1999 (see: http://www.oracle.com/database/documents/jserver_scalability_and_performance_twp.pdf). The JServer product available
10 from Oracle Corporation, USA, supports the concept of multiple sessions (a session effectively representing a transaction or application), each session including a JServer session. Resources such as read-only bytecode information are shared between the various sessions, but
15 each individual session appears to its JServer client to be a dedicated conventional JVM.

 US patent application 09/304160, filed 30 April 99 ("A long Running Reusable Extendible Virtual Machine"), assigned to IBM Corporation (IBM docket YOR9-1999-0170),
20 discloses a virtual machine (VM) having two types of heap, a private heap and a shared heap. The former is intended primarily for storing application classes, whilst the latter is intended primarily for storing system classes and, as its name implies, is accessible to
25 multiple VMs. A related idea is described in "Building a Java virtual machine for server applications: the JVM on OS/390" by Dillenberger et al, IBM Systems Journal, Vol 39/1, January 2000. Again this implementation uses a shared heap to share system and potentially application
30 classes for reuse by multiple workers, with each worker JVM also maintaining a private or local heap to store data private to that particular JVM process.

The above documents are focused primarily on the ability to easily run multiple JVMs in parallel. A different (and potentially complementary) approach is based on a serial rather than parallel configuration. Thus it is desirable to run repeated transactions (i.e. applications) on the same JVM, since this could avoid having to reload all the system classes at the start of each application. However, one difficulty with this is that each application expects to run on a fresh, clean, JVM. There is a danger with serial re-use of a JVM that the state left from a previous transaction somehow influences the outcome of a new transaction. This unpredictability is unacceptable in most circumstances.

US patent application 09/584641 filed 31 May 2000 in the name of IBM Corporation (IBM docket number GB9-2000-0061) discloses an approach for providing a JVM with a reset capability. US provisional application 60/208268 also filed 31 May 2000 in the name of IBM Corporation (IBM docket number YOR9-2000-0359) discloses the idea of having two heaps in a JVM. One of these is a transient heap, which is used to store transaction objects that will not persist into the next transaction, whilst a second heap is used for storing objects, such as system objects, that will persist. This approach provides the basis for an efficient reset mechanism by simply deleting the transient heap. The techniques described herein represent optimisations of the above methods, to allow the JVM reset to be performed as quickly and consistently as possible.

Summary of the Invention

Accordingly the invention provides a computer system providing an object-based environment, said computer system including storage, a contiguous linear portion of which is logically divided into first and second heaps located at opposite ends of the storage portion, with any gap between the two heaps representing an unallocated region of storage, wherein references are permitted from objects on the first heap to objects on the second heap and vice versa, said system comprising:

a garbage collector for operating across both heaps to remove objects that are no longer live;

means for expanding the first heap into said unallocated region according to a first expansion policy; and

means for expanding the second heap into said unallocated region according to a second expansion policy.

In a preferred embodiment, the computer system supports a transaction processing environment. The first heap is used for storing objects that are deleted at the end of the current transaction (and so can be regarded as a transient heap), and the second heap is used for storing objects that persist from one transaction to another. Since it is desired for each transaction to see as much as possible the same underlying system, the first heap is reset to the same predetermined initial size at the start of each transaction - an error condition is returned if the second heap has expanded such that it is not possible to reset the first heap to its predetermined initial size.

In the preferred embodiment, a midpoint is defined halfway between the first heap and second heap, when they each have their initial size (ie in the middle of the initially unallocated region). The first expansion policy is always to expand into the unallocated region in order to satisfy a storage request, whilst the second expansion policy is to expand into said unallocated region in order to satisfy a storage request until said midpoint is reached, whereupon said system preferentially performs a garbage collection to satisfy said request. Thus the first heap is generally freer to expand than the second heap. The reason for this is that the space will be reclaimed from the first heap at the end of the transaction, but this is not necessarily the case for the second heap. Nevertheless, the rate of expansion of the first heap into the unallocated region is slower once the first heap has passed said midpoint, in order to try to best conserve limited resources.

In the preferred embodiment, the second expansion policy further includes trying to shrink the first heap to allow room to expand said second heap in order to satisfy a storage request. In other words, if the first heap does not actually occupy all its allocated space, then it may be possible to transfer this to the second heap. This helps to optimize the use of storage between the two heaps.

Also in the preferred embodiment, the garbage collector optionally performs a compact operation after a garbage collection of the first and second heaps. The compact operation is performed in response to a first set of criteria relevant to the first heap, and a second set of criteria relevant to the second heap. Various complex

compaction policies can be adopted, but generally it is preferred that the second set of criteria are more sensitive to fragmentation than the first set of criteria. The reason for this is that the first heap can be deleted at the end of a transaction, and so its level of fragmentation is less of a concern than for the second heap which will persist for future transactions. It is preferred that after a compact operation, the first and second heaps are shrunk if possible, thereby returning released storage to said unallocated region. This helps to ensure that free space is available for future storage requests on either heap.

Although a single garbage collector is used for both heaps, the trigger for garbage collection and the type of garbage collection performed can vary between the first and second heaps, such as in accordance with their respective expansion policies. This allows the garbage collection strategy to be tuned to the requirements of each individual heap. This idea can be extended to only performing garbage collection on one heap at a time. For example, if a transaction is nearly finished, but more space is needed on the second heap, then it may not be necessary to also garbage collect the first heap, since this will soon be deleted. Likewise different garbage collection techniques (such as mark-sweep and copy garbage collection) can be used on different heaps, depending on their particular requirements.

In the preferred embodiment, the system further includes a bit array, having one bit for each possible object location in said portion of storage. The bit is used to provide a quick indication of whether or not there is an object currently stored at the corresponding

object location, which is useful in operations such as garbage collection

5 The invention further provides a method of operating a computer system providing an object-based environment, said computer system including storage, a contiguous linear portion of which is logically divided into first and second heaps located at opposite ends of the storage portion, with any gap between the two heaps representing an unallocated region of storage, wherein references are permitted from objects on the first heap to objects on the second heap and vice versa, said method comprising the steps of:

10 operating a garbage collector across both heaps to remove objects that are no longer live;

15 expanding the first heap into said unallocated region according to a first expansion policy; and

for expanding the second heap into said unallocated region according to a second expansion policy.

20 The invention further provides a computer program product comprising instructions encoded on a computer readable medium for causing a computer to perform the method described above. A suitable computer readable medium may be a DVD or computer disk, or the instructions may be encoded in a signal transmitted over a network from a server.

25 It will be appreciated that the method and computer program product of the invention will benefit from the same preferred features as the system of the invention.

Brief Description of the Drawings

A preferred embodiment of the invention will now be described in detail by way of example only with reference to the following drawings:

5 Figure 1 shows a schematic diagram of a computer system supporting a Java Virtual Machine (JVM);

 Figure 2 is a schematic diagram of the internal structure of the JVM;

10 Figure 3 is a flowchart depicting the steps required to load a class and prepare it for use;

 Figure 4 is a flowchart depicting at a high level the serial reuse of a JVM;

 Figure 5 is a schematic diagram showing the heap and its associated components in more detail;

15 Figures 6A and 6B form a flowchart illustrating garbage collection;

 Figure 7 is a flowchart illustrating heap expansion policy at a high level;

20 Figure 8 is a diagram of a lookup table used to determine if a reference is in a heap;

 Figure 9 is a diagram of a modified lookup structure for the same purpose as Figure 8, but for use in a system with much larger memory; and

25 Figures 10A and 10B form a flowchart illustrating the operations taken to delete the transient heap during JVM reset.

Detailed Description

30 Figure 1 illustrates a computer system 10 including a (micro)processor 20 which is used to run software loaded into memory 60. The software can be loaded into the memory by various means (not shown), for example from

a removable storage device such as a floppy disk, CD ROM, or DVD, or over a network such as a local area network (LAN), telephone/modem connection, or wireless link, typically via a hard disk drive (also not shown).

5 Computer system runs an operating system (OS) 30, on top of which is provided a Java virtual machine (JVM) 40. The JVM looks like an application to the (native) OS 30, but in fact functions itself as a virtual operating system, supporting Java application 50. A Java application may
10 include multiple threads, illustrated by threads T1 and T2 71, 72.

System 10 also supports middleware subsystem 45, for example a transaction processing environment such as CICS, available from IBM Corporation (CICS is a trademark of IBM Corporation). The middleware subsystem runs as an
15 application or environment on operating system 30, and initiates the JVM 40. The middleware also includes Java programming which acts to cause transactions as Java applications 50 to run on top of the JVM 40. In
20 accordance with the present invention, and as will be described in more detail below, the middleware can cause successive transactions to run on the same JVM. In a typical server environment, multiple JVMs may be running on computer system 10, in one or more middleware
25 environments.

It will be appreciated that computer system 10 can be a standard personal computer or workstation, network computer, minicomputer, mainframe, or any other suitable computing device, and will typically include many other
30 components (not shown) such as display screen, keyboard, sound card, network adapter card, etc which are not directly relevant to an understanding of the present

invention. Note that computer system 10 may also be an embedded system, such as a set top box, handheld device, or any other hardware device including a processor 20 and control software 30, 40.

5 Figure 2 shows the structure of JVM 40 in more detail (omitting some components which are not directly pertinent to an understanding of the present invention). The fundamental unit of a Java program is the class, and thus in order to run any application the JVM must first
10 load the classes forming and required by that application. For this purpose the JVM includes a hierarchy of class loaders 110, which conventionally includes three particular class loaders, named Application 120, Extension 125, and Primordial 130. An
15 application can add additional class loaders to the JVM (a class loader is itself effectively a Java program). In the preferred embodiment of the present invention, a fourth class loader is also supported, Middleware 124.

20 For each class included within or referenced by a program, the JVM effectively walks up the class loader hierarchy, going first to the Application class loader, then the Middleware loader, then the Extension class loader, and finally to the Primordial class loader, to see if any class loader has previously loaded the class.
25 If the response from all of the class loaders is negative, then the JVM walks back down the hierarchy, with the Primordial class loader first attempting to locate the class, by searching in the locations specified in its class path definition. If this is unsuccessful,
30 the Extension class loader then makes a similar attempt, if this fails the Middleware class loader tries. Finally, if this fails the Application class loader tries to load

the class from one of the locations specified in its class path (if this fails, or if there is some other problem such as a security violation, the system returns an error). It will be appreciated that a different class path can be defined for each class loader.

Note that if it is desired to load a further middleware class loader (i.e. one provided by the user rather than included within the JVM itself), then this can be achieved by declaring that the new class loader implements the middleware interface. This declaration by itself is sufficient for the JVM to treat it as a middleware class loader - no other method definitions or such-like are required.

The JVM further includes a component CL 204, which also represents a class loader unit, but at a lower level. In other words, this is the component that actually interacts with the operating system to perform the class loading on behalf of the different (Java) class loaders 110.

Also present in the JVM is a heap 140, which is used for storage of objects 145 (Figure 2 shows the heap 140 only at a high level; see Figure 5 below for more details). Each loaded class represents an object, and therefore can be found on the heap. In Java a class effectively defines a type of object, and this is then instantiated one or more times in order to utilise the object. Each such instance is itself an object which can be found in heap 140. Thus the objects 145 shown in the heap in Figure 2 may represent class objects or other object instances. (Note that strictly the class loaders as objects are also stored on heap 140, although for the sake of clarity they are shown separately in Figure 2).

Although heap 140 is shared between all threads, typically for reasons of operational efficiency, certain portions of heap 140 can be assigned to individual threads, effectively as a small region of local storage, which can be used in a similar fashion to a cache for that thread.

The JVM also includes a class storage area 160, which is used for storing information relating to the class files stored as objects in the heap 140. This area includes the method code region 164 for storing byte code for implementing class method calls, and a constant pool 162 for storing strings and other constants associated with a class. The class storage area also includes a field data region 170 for sharing static variables (static in this case implies belonging to the class rather than individual instances of the class, or, to put this another way, shared between all instances of a class), and an area 168 for storing static initialisation methods and other specialised methods (separate from the main method code 164). The class storage area further includes a method block area 172, which is used to store information relating to the code, such as invokers, and a pointer to the code, which may for example be in method code area 164, in JIT code area 185 (as described in more detail below), or loaded as native code such as C, for example as a dynamic link library (DLL).

Classes stored as objects 145 in the heap 140 contain a reference to their associated data such as method byte code etc in class storage area 160. They also contain a reference to the class loader which loaded them into the heap, plus other fields such as a flag (not

shown) to indicate whether or not they have been initialised.

Figure 2 further shows a monitor pool 142. This contains a set of locks (monitors) that are used to control access to an object by different threads. Thus when a thread requires exclusive access to an object, it first obtains ownership of its corresponding monitor. Each monitor can maintain a queue of threads waiting for access to any particular object. Hash table 141 is used to map from an object in the heap to its associated monitor.

Another component of the JVM is the interpreter 156, which is responsible for reading in Java byte code from loaded classes, and converting this into machine instructions for the relevant platform. From the perspective of a Java application, the interpreter effectively simulates the operation of a processor for the virtual machine.

Also included within the JVM are class loader cache 180 and garbage collection (GC) unit 175. The former is effectively a table used to allow a class loader to trace those classes which it initially loaded into the JVM. The class loader cache therefore allows each class loader to check whether it has loaded a particular class - part of the operation of walking the class loader hierarchy described above. Note also that it is part of the overall security policy of the JVM that classes will typically have different levels of permission within the system based on the identity of the class loader by which they were originally loaded.

Garbage collection (GC) facility 175 is used to delete objects from heap 140 when those objects are no

longer required. Thus in the Java programming language, applications do not need to specifically request or release memory, rather this is controlled by the JVM. Therefore, when Java application 50 creates an object 145, the JVM secures the requisite memory resource. Then, when Java application 50 finishes using object 145, the JVM can delete the object to free up this memory resource. This latter process is known as garbage collection, and is generally performed by briefly interrupting all threads 71, 72, and scanning the heap 140 for objects which are no longer referenced, and hence can be deleted. The garbage collection of the preferred embodiment is described in more detail below.

The JVM further includes a just-in-time (JIT) compiler 190. This forms machine code to run directly on the native platform by a compilation process from the class files. The machine code is created typically when the application program is started up or when some other usage criterion is met, and is then stored for future use. This improves run-time performance by avoiding the need for this code to be interpreted later by the interpreter 156.

Another component of the JVM is the stack area 195, which is used for storing the stacks 196, 198 associated with the execution of different threads on the JVM. Note that because the system libraries and indeed parts of the JVM itself are written in Java, and these frequently use multi-threading, the JVM may be supporting multiple threads even if the user application 50 running on top of the JVM contains only a single thread itself.

It will be appreciated of course that Figure 2 is simplified, and essentially shows only those components

pertinent to an understanding of the present invention. Thus for example the heap may contain thousands of Java objects in order to run Java application 50, and the JVM contains many other components (not shown) such as diagnostic facilities, etc.

Figure 3 is a flowchart illustrating the operations conventionally performed to load a class in order to run a Java application. The first operation is loading (step 310) in which the various class loaders try to retrieve and load a particular class. The next operation is linking, which comprises three separate steps. The first of these is verification (step 320), which essentially checks that the code represents valid Java programming, for example that each instruction has a valid operational code, and that each branch instruction goes to the beginning of another instruction (rather than the middle of an instruction). This is followed by preparation (step 330) which amongst other things creates the static fields for a class. The linking process is completed by the step of resolution, in which a symbolic reference to another class is typically replaced by a direct reference (step 340).

At resolution the JVM may also try to load additional classes associated with the current class. For example, if the current class calls a method in a second class then the second class may be loaded now. Likewise, if the current class inherits from a superclass, then the superclass may also be loaded now. This can then be pursued recursively; in other words, if the second class calls methods in further classes, or has one or more superclasses, these too may now be loaded. Note that it is up to the JVM implementation how many classes are

loaded at this stage, as opposed to waiting until such classes are actually needed before loading them.

The final step in Figure 3 is the initialisation of a loaded class (step 350), which represents calling the static initialisation method (or methods) of the class. According to the formal JVM specification, this initialisation must be performed once and only once before the first active use of a class, and includes things such as setting static (class) variables to their initial values (see the above-mentioned book by Lindholm and Yellin for a definition of "first active use"). Note that initialisation of an object also requires initialisation of its superclasses, and so this may involve recursion up a superclass tree in a similar manner to that described for resolution. The initialisation flag in a class object 145 is set as part of the initialisation process, thereby ensuring that the class initialisation is not subsequently re-run.

The end result of the processing of Figure 3 is that a class has been loaded into a consistent and predictable state, and is now available to interact with other classes. In fact, typically at start up of a Java program and its concomitant JVM, some 1000 objects are loaded prior to actual running of the Java program itself, these being created from many different classes. This gives some idea of the initial delay and overhead involved in beginning a Java application.

As mentioned above, the problems caused by this initial delay can be greatly reduced by serial reuse of a JVM, thereby avoiding the need to reload system classes and so on. Figure 4 provides a high-level flowchart of a preferred method for achieving such serial reuse. The

method commences with the start of the middleware subsystem 45, which in turn uses the Java Native Interface (JNI) to perform a Create JVM operation (step 410). Next an application or transaction to run on the JVM is loaded by the Application class loader 120. The middleware includes Java routines to provide various services to the application, and these are also loaded at this point, by the Middleware class loader 124.

The application can now be run (step 420), and in due course will finally terminate. At this point, instead of terminating the JVM as well as the application, the middleware subsystem makes a Reset JVM call to the JVM (step 430). The middleware classes may optionally include a tidy-up method and/or a reinitialize method. Both of these are static methods. The JVM responds to the Reset JVM by calling the tidy-up method of the middleware classes (step 440). The purpose of this is to allow the middleware to leave the JVM in a tidy state, for example removing resources and closing files that are no longer required, and deleting references to the application objects. In particular, all those middleware classes which have been used since the previous JVM reset (or since the JVM was created if no resets have occurred) have their tidy-up method called, assuming of course that they have a tidy-up method (there is no requirement for them to have such a tidy-up method).

The tidy-up method may be similar to the finalise method of a class, which is a standard Java facility to allow an object to perform some close-down operation. However, there is an important difference in that tidy-up is a static method. This means that contrary to the finalise method it applies to the class rather than any

particular object instance, and so will be called even if there are no current object instances for that class. In addition the timing of the tidy-up method is different from finalise, in that the former is called in response to a predetermined command to reset the JVM. In contrast, in accordance with the JVM specification, the finalise method is only triggered by a garbage collection. More particularly, if an object with a finalizer method is found to be unreachable during a garbage collection (ie it is no longer effectively active) then it is queued to the finalizer thread, which then runs the finalizer method after the garbage collection is completed. Note that the finalizer method of an object may never be called, if an application finishes and the JVM shuts down without the system needing to perform a garbage collection.

Once the tidy-up has been completed, a refresh heap operation is performed (step 445). As will be described in more detail below, this deletes those portions of the heap that relate to the application or transaction that has just been completed, generally analogous to a garbage collection cycle. Note that many of the objects deleted here might not have been removable prior to the tidy-up method, since they could still have been referenced by the middleware classes.

At this point, the middleware subsystem makes a determination of whether or not there is another application to run on the JVM (step 450). If not, the middleware subsystem uses the JNI to make a Destroy JVM call (step 460) which terminates the JVM, thereby ending the method of Figure 4. If on the other hand there is another application to run, then this new application is

started by the middleware. The system responds to this new application by calling in due course the reinitialisation method in each of the middleware classes to be reused (step 455). The purpose of this is to allow the middleware classes to perform certain operations which they might do at initialisation, thereby sidestepping the restriction that the JVM specification prevents the initialisation method itself being called more than once. As a simple example, the reinitialisation may be used to reset a clock or a counter. As shown in Figure 4, the system is now in a position to loop round and run another application (step 420).

It is generally expected that the reinitialisation method will be similar in function to the initialisation method, but there may well be some differences. For example, it may be desired to reset static variables which were initialised implicitly. Another possibility is to allow some state or resources to persist between applications; for example, if a class always outputs to one particular log file which is set up by the initialisation method, it may be more efficient to keep this open in between successive JVMs, transparent to the application.

It should be noted that whilst Figure 4 indicates the distinct logical steps performed by the method of the invention, in practice these steps are not all independent. For example, calling the tidy-up methods (step 440) is part of the overall reset JVM operation (step 430). Likewise, calling the reinitialisation methods (step 455) is effectively part of the start-up processing of running the new application (step 420). Thus reinitialisation is performed prior to first active

use of a class, and this may occur at any stage of a program. Therefore class reinitialisation (like conventional initialisation) is not necessarily completed at start-up of the program, but rather can be regarded as potentially an ongoing process throughout the running of a program.

It will also be appreciated that there is some flexibility with regard to the ordering of the steps shown in Figure 4. In particular, the decision of whether or not there is to be another application (step 450) could be performed earlier, such as prior to the refresh heap step, the tidyup step, and/or the reset JVM step. In the latter case, which corresponds to immediately after the first application has concluded (i.e. straight after step 420), the alternative outcomes would be to destroy the JVM (step 460) if there were no further applications, or else to reset the JVM, tidy up, refresh the heap, and reinitialise (steps 430, 440, 445, and 455) if there were further applications. If instead the decision step 450 is intermediate these above two extreme positions, the logic flow can be determined accordingly. Further details about the implementation of the tidyup and reinitialise methods are provided in above-mentioned US patent application 09/584641.

It should be noted that in the preferred embodiment, the ability to reset the JVM, and to have tidyup and reinitialise methods, is only available for middleware classes (i.e. those loaded by the middleware class loader). This is to allow the middleware classes to be re-used by successive applications or transactions, for which they can perform various services. The basis for this approach is that typically the middleware is a

relatively sophisticated and trusted application, and so can be allowed to take responsibility for proper implementation of the tidy-up and reinitialise methods. On the other hand, the transactions that run within the
5 middleware are not treated as reliable. Note also that the system classes themselves do not have tidyup or reinitialisation methods, despite persisting across a JVM reset. Rather, if the middleware makes any change to a system class, then the middleware itself is expected to
10 take the necessary action (if any) for a reset with respect to the system class as part of the middleware's own tidyup operation.

An important part of the reset JVM/tidyup operation (steps 430 and 440) in the preferred embodiment is to
15 make sure that the JVM is in a state which is amenable to being tidied up. If this is the case, the JVM is regarded as being clean, if not, it is regarded as being dirty or contaminated.

Considering this in more detail, if the application
20 has performed certain operations, then it will not be possible for the middleware classes to be certain that their tidy-up and reinitialise methods will fully reset the system to a fresh state. With such a contaminated JVM, the system still calls the tidy-up methods of the
25 class objects as per normal (step 440), but the return code back to the middleware associated with the reset JVM operation (step 430) effectively indicates failure. The expectation here is that the JVM would actually be
30 terminated by the middleware subsystem at this point, as it is no longer in a predictable condition.

One important situation which would prevent the JVM from being able to properly reset is where the

application has performed certain operations directly such as making security or environment changes, running native code, or performing Abstract Windowing Toolkit (AWT) operations. These affect the state of the JVM or the underlying computer system and cannot be reliably tidied up by the middleware, for the simple reason that the middleware does not necessarily know about them. Such changes could then persist through a reset JVM call, and contaminate the JVM for any future applications. In contrast, if an application performs such operations through a middleware call, then this does not cause any problems, because the middleware now does know about the situation and so can perform whatever tidyup measures are required.

The JVM thus monitors for operations that may prevent proper reset, including whether they have been performed by an application or middleware. This is determined by the JVM keeping track of its context, which is set to application context for an application class, and to middleware context for a middleware class, whilst a primordial or extension class has no impact on the existing context of application or middleware. In particular, context can be determined based on the type of class which contains the method that is currently being performed, whilst the type of class is determined from its original class loader.

As previously mentioned, the list of problematic operations given above only causes difficulty when performed in an application context, since in a middleware context it is possible for them to be reset by the appropriate tidy-up routines of the relevant middleware classes.

Referring now to Figure 5, in the preferred embodiment the heap 140 is logically split into three components (objects in one component can reference objects in another component). In particular, at the bottom (logically) of heap 140 is middleware section 510, and at the top of the heap is transient section 520. The data in these two heaps grows towards each other, thus transient heap grows in the direction of arrow 521, and middleware heap in the direction of arrow 511. The middleware heap is defined by boundary 512, and the transient heap by boundary 522, with unassigned space 515 between them. It should be appreciated that boundaries 512 and 522 represent the maximum size currently assigned to the two heaps, rather than their current fill levels - these are instead shown by dashed lines 513 and 523. In other words, as the middleware heap fills up, the fill level 513 will approach towards middleware heap boundary 512; likewise as the transient heap fills up, the fill level 523 will approach towards transient heap boundary 522. Finally, and separate from the transient heap and middleware heap, is system heap 550. Note that the combined transient and middleware heaps, together with intervening unassigned space, are allocated from a single physically contiguous block of memory 560. In contrast, the system heap 550 may be formed from multiple non-contiguous regions of memory.

In one preferred embodiment, memory 560 comprises 64 MBytes, and the initial size of the middleware and transient heaps is 0.5 Mbyte each. Thus it can be seen that initially the unassigned region 515 dominates, although as will be discussed in detail below, the transient and middleware heaps are allowed to expand into

this space. However, these values are exemplary only, and suitable values will vary widely according to machine architecture and size, and also the type of application.

5 Heap control block 530 is used for storing various information about the heap, such as the location of the heap within memory, and the limits of the transient and middleware sections as defined by limits 512 and 522. Free chain block 532 is used for listing available storage locations within the middleware and transient
10 sections (there is actually one free chain block for each section). Thus although the middleware and transient heaps start to fill sequentially, the likely result of a garbage collection cycle is that space may become available within a previously occupied region. Typically
15 therefore there is no single fill line such as 513, 523 between vacant and occupied space, rather a fragmented pattern. The free chain block is a linked list which specifies the location and size of empty regions within that section of the heap. It is quick to determine
20 whether and where a requested amount of storage is available in the heap by simply scanning through the linked list. Note that in the preferred embodiment, empty regions in the heap which are below a predetermined size (typically a few hundred bytes) are excluded from the
25 free chain list. This prevents the list from becoming too long through containing a large number of very small vacant regions, although it does mean that these regions effectively become inaccessible for storage (although they can be retrieved later, as described in more detail
30 below).

 The transient heap 520 is used for storing objects having no expected currency beyond the end of the

application or transaction, including application object instances, and primordial object instances and arrays created by application methods (arrays can be regarded as a specialised form of object). Since the lifetime of such objects is commensurate with the application itself, it should be possible to delete all the objects in the transient heap at the end of the application. The application class objects are also on the transient heap. In contrast, the middleware heap 510 is used for storing objects which have a life expectancy longer than a single transaction, including middleware object instances, and primordial object instances and arrays created by middleware methods. In addition, string objects and arrays for strings interned in the Interned String Table are also stored in the middleware heap (the Interned String Table is a tool whereby if multiple identical strings are to be stored on the heap, it is possible to store only one copy of the string itself, which can then be referenced elsewhere). Lastly, the system heap 550 is used for storing primordial class objects and reusable class objects, where the term reusable class object is used to denote a class which can be used again after JVM reset.

The type of class is dependent on the class loader which originally loaded it, in other words a middleware class and an application class are loaded by the middleware class loader 124 and the application class loader 120 respectively. For the purposes of the present discussion, primordial classes can be considered as classes loaded by the Primordial or Extensions class loader (130 and 125 respectively in Figure 2). In the

preferred embodiment, classes loaded by the middleware class loader are automatically regarded as reusable.

It is clear from above that instances of primordial classes, such as the basic string class java/lang/String, can be located either in the middleware heap or the transient heap, depending on the method which created them. In a preferred embodiment of the present invention, the determination of where to place such primordial class instances is based on the current context described above (also referred to as method-type). Thus if a method belonging to an application class is invoked, the context or method-type becomes Application, whilst if a method belonging to a middleware class is invoked, the method-type becomes Middleware. Finally, if a method belonging to a primordial class is invoked, the method-type is unchanged from its previous value. The context or method-type is stored in the Java frame for the method (which is stored on stack 195 - see Figure 2); at the completion of the method, the method-type reverts to its value at the time the method was invoked, which was stored in the previous frame.

It should be noted that for the above purpose a method belongs to the class that actually defines it. For example, if class A subclasses class B, but does not override method C, then method C belongs to class B. Therefore the method-type is that of class B, even if method C is being run for an instance of class A. In addition, the reason for tracking method-type on a per-thread basis is that it is possible for various threads within an application to be executing different methods having different context.

The transient region of the heap, containing objects created by the application or transaction, is subject to normal garbage collection, but the intention is that it will be sufficiently large that this is unlikely to occur within the lifetime of a typical application. At the end of each application, the transient region of the heap is reset. (The repetition of this pattern will thereby avoid having to perform garbage collection during most typical applications). In contrast the middleware region generally contains objects created by the trusted middleware. It is again subject to conventional garbage collection, although in a transaction environment it is expected that the majority of objects will be created in the transient heap, so that garbage collection is not expected to occur frequently. Moreover the system typically tries to perform garbage collection of the middleware heap at the same time as reset of the transient heap, in other words between rather than during transactions (this is discussed in more detail below). The middleware heap is not cleared between applications, but rather remains to give the middleware access to its persistent state (it is assumed that the middleware can take responsibility for resetting itself to the correct state to run the next application).

The preferred embodiment is actually somewhat more complicated than described above, in that it supports two types of application class loader, one of which is for standard application classes, the other for reusable application classes. The motivation here is that when the next transaction is to run, it will in fact require many of the same application classes as the previous transaction. Therefore it is desirable to retain some

application system classes rather than having to reload them, although certain additional processing is required to make them look newly loaded to the next transaction. Conversely it would be possible to have a second
5 middleware class loader which is for non-reusable middleware classes. In the former situation the reusable application classes are treated essentially in the same manner as the reusable middleware classes, (eg loaded into the system heap); in the latter situation the
10 non-reusable middleware classes would be treated similarly to the non-reusable application classes but loaded into the middleware heap (since they may exist after the conclusion of a transaction, even if they do not endure for the next transaction). However, for
15 present purposes in order to explain the invention more clearly, it will be assumed that all the middleware classes are reusable, and that none of the application classes are reusable.

The introduction of multiple heaps for different
20 types of objects allows the handling of the heap to be fine-tuned to the requirements of those types of object. For example, it may be desirable for the transient heap to allocate a larger thread local heap cache. In addition, utilising a single block of memory for the
25 transient and middleware heaps improves space usage, in that a given region of memory can be flexibly assigned to either the transient or middleware heap, depending on particular application requirements. On the other hand it does lead to some complications in terms of heap
30 management, especially as regards control of heap size. Thus in simple terms, as more and more objects are created, there is a choice to either enlarge the size of

the heap, or to perform a garbage collection to maintain the heap within current size limits. The former option is generally quick, but will eventually lead to the exhaustion of heap space; in contrast, a garbage collection is relatively slow, since it interrupts processing, but does constrain the heap size to within predetermined limits. Overall, the preferred embodiment tries to avoid garbage collections during transactions as much as possible, thereby optimising performance for the transaction, and to rely instead on the heap refresh described below, which is performed at the end of the transaction as part of the JVM reset.

More specifically, the policy for expansion and garbage collection in terms of system heap 550 is straightforward, in that objects in this heap are never garbage collected; rather this heap simply expands to accommodate all relevant class objects. However, the policy for transient and middleware heaps is more complex, because these two heaps are interdependent, in that they share the same memory space. In order to better understand this policy, it will be helpful to firstly review in more detail the garbage collection strategy of the preferred embodiment, as shown in Figures 6A and 6B. In particular, the method involves firstly a mark phase, which marks all objects in the heap that are currently in use (known as live or active objects), and secondly a sweep phase, which represents the actual deletion of objects from the heap. Note that general background on garbage collection algorithms can be found in "Garbage Collection: Algorithms for Automatic Dynamic Memory Management" by R Jones and R Lins, Wiley, 1996 (ISBN 0 471 94148 4), whilst one implementation for garbage

collection in a system having multiple heaps is described in: "A customisable memory management framework for C++" by G Attardi, T Flagella, and P Iglío, in Software Practice and Experience, vol 28/11, 1998.

5 As shown in Figure 6A, the method starts with a review of the registers and stack, both the Java stack, as shown in Figure 2, and also the C stack, (assuming that the JVM 40 is running as a C application on OS 30, see Figure 1) (step 610). Each thirty-two bit data word
10 (for a 32-bit system) contained therein could represent anything, for example a real number, or part of a string, but it is assumed at least initially that it may denote a 32 bit reference to an object location in the heap. To firm up on this assumption, three tests are made.
15 Firstly, it is tested whether or not the number references a location within the heap (step 612); if not then the number cannot represent an object reference. Secondly, in the preferred embodiment, all objects commence on an 8-byte boundary. Thus if the location
20 corresponding to the data word from the stack/register does not fall on an object boundary (tested at step 615), then the original assumption that the data/number represents a reference to the heap must again be rejected. Thirdly, in the preferred embodiment, a table
25 538 is maintained (see Figure 5) which has a bit for each object location in the heap; this bit is set to unity if there is an object stored at that location, and zero if no object is stored at that location (the relevant bit is updated appropriately whenever an object is created,
30 deleted, or moved). If the data word from the stack/register corresponds to an object location for which the bit is zero, in other words, no object at that

location, then once more the original assumption that the data/number represents a reference to the heap must be rejected (step 620). If the data word passes all three of the tests of steps 612, 615 and 620, then there are three remaining possibilities: (a) the word references an object on the heap; (b) the word is an integer that happens to have the same value as the object reference; or (c) the word is a previous value from uninitialized storage. As a conservative measure, it is assumed that option (a) is correct, and so the object is marked as live (step 625). A special array of bits is provided (block 534, see Figure 5), one bit per object, in order to store these mark bits. If there remain other values on the stacks/registers to test (step 630), the method then loops back to examine these in the same manner as just described; if not the first stage of the mark process is complete.

In the second stage of the mark process, shown in Figure 6B, the objects marked as live are copied onto a list of active objects (step 635) (in the preferred embodiment objects are actually copied to the active list when originally marked, ie at the same time as step 625 in Figure 6A). An object from this list is then selected (step 640), and examined to see if it contains any references (step 645). Note that this is a reasonably straightforward procedure, because the structure of the object is known from its corresponding class file, which defines the relevant variables to be used by the object. Any objects referenced by the selected object are themselves marked (step 650) and added to the active list (step 655). Next, the selected object is removed from the active list (step 660), and then a test is performed

(step 665) to determine if the active list is empty; if not, processing loops back to step 640 to select another object from the active list. Finally, when step 665 produces a positive outcome, all objects that are active, because they are referenced directly or indirectly from the stacks or registers, have been appropriately marked.

The mark stage is then followed by a sweep stage (step 670) and a compact stage (step 675). The former garbage collects (ie deletes) all those objects which have not been marked, on the basis that they are no longer reachable from any live or active object. In particular, each object which is not marked as active has its corresponding bit set to zero in table 538 (see Figure 5). Runs of zeros in the bit allocation table 538 are now identified; these correspond to some combination of the object immediately preceding the run, which may extend into the run (since only the head of an object is marked in the bit allocation table), and free space (released or never filled). The amount of free space in the run of zeros can be determined by examining the size of the object immediately preceding the run. If the amount of free space exceeds the predetermined minimum amount mentioned earlier, then the run is added to the free chain list 532 (see Figure 5).

Over time, such sweeping will tend to produce many discontinuous vacant regions within the heap, corresponding to the pattern of deleted objects. This does not represent a particularly efficient configuration, and in addition there will be effective loss of those pieces of memory too small to be on the free list. Hence a compact stage (step 675) can be performed, which acts to squeeze together those objects

which remain in the heap after the sweep in order to
amass them into a single continuous block of storage (one
for the transient heap, one for the middleware heap).
Essentially, this means relocating objects from their
initial positions in the heap, to a new position so that,
as much as possible, they are all adjacent to one
another. As part of this compaction, the very small
regions of memory too small to be on the free chain 532
(see Figure 5) should be aggregated into larger blocks
that can be recorded in the free chain.

An important requirement of the object relocation of
the compaction step is of course that references to a
moved object are altered to point to its new location.
This is a relatively straightforward operation for object
references on the heap itself, since as previously
mentioned, they can be identified from the known
structure of each object, and updated to the appropriate
new value. However, there is a problem with objects which
are directly referenced from a register or stack. As
discussed above, each number in the register/stack is
treated for garbage collection purposes as if it were an
object reference, but there is no certainty that this is
actually the case; rather the number may represent an
integer, a real number, or any other piece of data. It is
therefore not possible to update any object references on
the stack or register, because they may not in fact be an
object reference, but rather some other piece of program
data, which cannot of course be changed arbitrarily. The
consequence of this is that it is impossible to move an
object which appears to be directly referenced from the
heap or stack; instead these objects must remain in their
existing position. Such objects are informally known as

"dozed" objects since they cannot be moved from their current position.

Two other classes of objects which cannot be moved from the heap are class objects, and thread objects (thread objects are control blocks used to store information about a thread). The reason for this is that such objects are referenced from so many other places in the system that it is not feasible to change all these other references. These objects are therefore known as "pinned", since like dozed objects they cannot be moved from their current position.

A consequence of pinned and dozed objects is that a compact process may not be able to accumulate all objects in a heap into a single contiguous region of storage, in that pinned and dozed objects must remain in their original positions. The consequences of this are discussed in more detail below.

Note that in the preferred embodiment, a compact stage (step 675) is not necessarily employed on every garbage collection cycle, unless this is explicitly requested as a user initial set-up option. Rather a compact operation is only performed when certain predetermined criteria are met. For example, as previously indicated a garbage collection can be triggered by a request for storage in the heap that cannot be satisfied. If the request still cannot be satisfied after the sweep step 670, because there is no single block of memory available of sufficient size, then a compact stage is automatically performed, to try and accumulate an adequately-sized storage region.

In the preferred embodiment, the further criteria used for deciding whether to compact are different for

the middleware heap and the transient heap. Thus for the transient heap a compaction is performed whenever the amount of free space remaining in the transient heap after the garbage collection is less than 5% of the heap capacity. The idea here is that when space appears to be running out, the compacting should retrieve some additional space from those empty regions too small for the free chain list. On the other hand, for the middleware heap more complex compaction algorithms are used, based for example on when heap fragmentation exceeds certain limits (eg in terms of number of fragments), or where the largest block in the free chain list is below a certain size. The rationale here is that the middleware heap is likely to be of relatively long duration, and so it is worthwhile to try to optimise its overall storage arrangement.

Note that although the triggers for garbage collection and compaction can be different for the middleware and transient heap, when either operation is performed, in the preferred embodiment it is performed on the whole of active storage 560 - ie on both the middleware and transient sections simultaneously. This is because interheap references are permitted, and so any marking or compaction operation necessarily involves both heaps. Consequently, once starting a garbage collection or compaction, it is most efficient to do both heaps at the same time.

One complication to the garbage collection described above is that as previously mentioned, Java permits objects to have finalizer methods, which must be run prior to deletion of the object in a garbage collection. In order to manage this requirement, certain additional

processing is required (not shown in Figure 6). Thus when an object is created on the heap that has a finalizer method, a reference to that object is added to a set of finalizer references. At the end of the mark phase of garbage collection, this set of finalizer references is scanned, to detect any objects in the set which are not marked - the resultant group represents the objects which are about to be deleted, and so need to have their finalizer methods run. To accomplish this, objects in this group now need to be marked as live, and their references iteratively traced and also marked as live, in similar fashion as for the main mark phase. The purpose of this is firstly to retain the objects in order to run their finalizer methods, and secondly to retain any other objects which are directly or indirectly referenced by them, so that the finalizer methods run correctly. The finalizer references for objects in this group are removed from the set of finalizer references described above, so that their finalizer method will not be activated by any future garbage collection cycle, and passed to a reference handler. The subsequent processing is asynchronous, and does not occur until main system processing is resumed after the garbage collection has concluded (ie after the end of the processing of Figure 6B). Once the reference handler has restarted, it passes any object finalizer references it received during the garbage collection to a finalizer queue. A separate finalizer thread then runs each entry in the queue in turn, deleting the object reference from the queue after the corresponding finalizer method has been run.

Note that objects referenced by the reference handler or on the finalizer queue are regarded as "live"

during a garbage collection process. In other words they are marked along with any other objects which they reference, directly or indirectly. This ensures that objects do not get inadvertently deleted from the finalizer queue, if their wait on this queue exceeds the time to the next garbage collection. (Thus objects in the reference handler and finalizer queue form additional roots for live objects, in addition to those on the stacks and registers as illustrated in Figure 6; in fact in the preferred embodiment, there are other categories of roots, for example system class files, but the details are not pertinent to an understanding of the present invention).

One potential problem with the handling of finalizer methods described above is that by running them on a dedicated thread (the finalizer thread), the context of the thread will be different from the main application thread, where context here indicates general system properties associated with the thread, such as security permissions. This can be a particular concern in relation to transaction threads, which as previously mentioned are regarded as relatively untrustworthy. Therefore, the preferred embodiment modifies the handling of objects in the transient heap having finalizer methods. If these are located in a garbage collection cycle and are not marked, then as described above they are marked, along with the objects which they reference, directly or indirectly. However, no further processing is done on these objects, in particular, they are not removed from the set of finalizer references, and are not passed to the reference handler. The effect of this is that these object then simply continue to appear to the garbage collection

process as normal live objects, and are maintained through each garbage collection cycle. These objects are eventually deleted in the refresh heap step 445 of the JVM reset (see Figure 4), which will be described in more detail below.

Returning now to the question of allocating heap space from the overall memory region 560, which contains both the middleware and transient sections, the procedure for this is illustrated at a high level in Figure 7 (at this level the same general policy is used for both the middleware and transient heaps, although as will be seen below, there are some significant differences in the details of their respective policies). The process starts with an allocation request (step 705), typically to store an object on the heap. This causes the free chain block 532 (see Figure 5) for the relevant heap section to be examined; if there is available space (step 715), then the method proceeds directly to allocating the desired space (step 795), and exits successfully.

On the other hand, if the test of step 715 is negative, then it means that the heap is too full to sustain the new allocation. This is equivalent conceptually to the fill level 513 in Figure 5 approaching the assigned boundary 512 for the middleware heap, or fill level 523 approaching assigned boundary 522 for the transient heap. In this situation, the system first determines whether it is possible to simply expand the amount of space assigned to the heap (step 725). In simple terms, for the middleware heap this corresponds to moving assigned boundary 512 upwards into the unassigned region 515, thereby taking some of the unassigned storage and allocating it to the middleware heap 510; conversely

for the transient heap, boundary 522 is moved downwards. Of course, it is not possible for the middleware heap to encroach into the transient heap or vice versa, so that once the unassigned space 515 has been exhausted, then it is no longer possible to expand the heaps further. In a situation where heap space is available, then a policy is defined to determine the amount of extra space to add to the heap. The general policy in the preferred embodiment is to increase the heap so that there is 30% free space (taking into account the new allocation request).

However, a predetermined minimum expansion size is defined (0.5 MByte in the preferred embodiment), so that the expansion is actually 30% or 0.5 MByte, whichever is greater (subject of course to the amount of space available). Likewise, the user may also set a maximum expansion size, which is then used to cap the figure just obtained (providing it does not prevent satisfying the current allocation request). Finally, in the preferred embodiment, heap memory is always assigned/deassigned in units of a predetermined size, which for a 32-bit system is 64 Kbytes for reasons that will be described later. Therefore whatever expansion value is determined based on the 30% expansion, this is adjusted to the appropriate whole number of 64 Kbyte units. Note that in the preferred embodiment, there are further controls on how the different heaps are allowed to expand; these are discussed below in more detail.

After the available expansion space has been determined, it is tested whether there will now be sufficient space to satisfy the allocation request (step 735). If so, the relevant heap is duly expanded (step 785), if not, the method proceeds to step 745, and a

garbage collection is performed. It is now checked whether or not this has created sufficient space (step 755); if so, the method proceeds to allocate the requested space (step 795). Note that one minor complexity not shown in Figure 7 is that the garbage collection (step 745) may perform both a compact operation, and then also try a heap expansion (equivalent to step 785), if these are necessary to obtain the requested space. If on the other hand there is still insufficient space for the allocation request, then as a final measure, it is possible to shrink the other heap (step 765). Thus referring back to Figure 5, it can be seen that middleware heap could in principle lose the assigned but empty space between boundary 512 and fill level 513, by lowering boundary to fill level 513. The reclaimed space could then be transferred to the transient heap 520 (assuming that it already now extended through the region 515 shown in Figure 5 as unassigned). Conversely, space could be made available for transfer from the transient heap to the middleware heap by raising boundary 522 towards fill level 523.

Following the shrinkage of the other heap (step 765) a test is now made to see if this has created sufficient space for the allocation request (step 775); if not the system must return an error to the allocation request (step 780) indicating that no space is available. Assuming however that space is available, then the heap for which the allocation request is made can expand (step 785) into the space vacated by the shrinkage of the other heap, thereby allowing the allocation request to be satisfied (step 795).

It will be appreciated that there are many possible variations on the processing shown in Figure 7. For example Figure 7 shows heap expansion (step 785) only when this will positively provide the required space (ie following a positive result from the tests of steps 735, 765, 775), but it will be appreciated that such heap expansion might be performed irrespective of whether or not this would create sufficient space for the allocation request for some or all of these tests. In fact, in the preferred embodiment, after garbage collection has been performed (step 745), the relevant heap will automatically try to expand to give 30% free space as previously described, even when the allocation request has already been satisfied (this is subject to certain limitations described in more detail below).

In addition, an attempt could be made to shrink the other heap (step 765) before performing garbage collection (step 745), or it may occur automatically as part of the garbage collection process. Thus in the preferred embodiment, the assigned boundary for the transient heap (line 522 in Figure 5) is shrunk as much as possible each time the heap has been compacted, providing that this does not reduce the transient heap below its initial size. In contrast, although the middleware heap is also shrunk after compaction in the preferred embodiment, in general some leeway (such as 30% free space) is left between the heap boundary and the fill level. The middleware heap is also never reduced below its original size. This policy balances the fact that the transient heap is allowed to grow more easily than the middleware heap (as discussed below). More generally, such shrinkage after compaction returns

storage to the unassigned pool, and so increases flexibility for managing storage requests from the two heaps. Note that because in the preferred embodiment shrinkage is performed (if possible) after compaction, which in turn will be performed if the garbage collection does not otherwise satisfy the allocation request, then to some extent steps 745 and 765 in Figure 7 are effectively amalgamated together.

Although the processing shown in Figure 7 applies at a high level, there are important differences in detail as regards the management of the transient and middleware heaps. The policies adopted reference a location 565 which represents the midpoint between the middleware heap boundary 512 and the transient heap boundary 522 (see Figure 5), as determined at JVM start-up or JVM reset. Thus for the middleware heap, the procedure is expand the heap rather than garbage collect, using the expansion criteria described above, until the heap would expand past the midpoint location 565. If this situation does arise, then the system uses a smaller expansion increment, namely the minimum expansion value (ie 0.5 Mbyte in the preferred embodiment). Finally, if even this reduced expansion would still take the middleware heap past the midpoint, then a garbage collection is performed (ie step 745), rather than allowing the middleware heap to expand further. As previously indicated, a compaction will be performed here if necessary to satisfy the allocation request. After the garbage collection, the system will then try to expand the middleware heap using the standard policy based on 30% free space, or the minimum expansion value of 0.5 Mbyte if the 30% expansion would exceed the midpoint. In other words, the policy is

to try to prevent the middleware heap from expanding past midpoint 565 (although this may happen eventually if the garbage collection does not reclaim sufficient space). The rationale behind this is to try to avoid taking up space from the transient heap, a particular concern being the possibility of a long-lived middleware object becoming pinned high up (in the sense of Figure 5) in the heap storage 560, and therefore seriously limiting the amount of space available to the transient heap.

Considering now the transient heap, then once this reaches (or would reach) the midpoint 565, then again the expansion rate for this heap is reduced to half the minimum expansion value. However, unlike for the middleware heap, this expansion is allowed to continue on past the midpoint, until eventually all usable heap space is exhausted, when clearly a garbage collection will be needed. The motivation here is that it is expected that most new objects for the transaction will be created on the transient heap, so that this requires most room. Moreover, since the transient heap will be deleted anyway at the conclusion of the transaction, the concern about pinned objects is reduced (or the JVM will become dirty, as discussed in more detail below). A further consequence of this is that there is a general desire for performance reasons if possible to avoid a garbage collection during a transaction, but rather to postpone this if possible until the heap refresh (step 445, see Figure 4) performed as part of the JVM reset.

With reference to step 765 in Figure 7 (shrinking the other heap to reclaim space), this step is not performed for an allocation request to the transient heap (in other words, a No from step 755 would go straight to

Error 780). However, it will be noted that if the allocation request is about to fail, the heap would already have been garbage collected and compacted, and the size of the heaps shrunk as per the policy discussed above, so that the amount of free space available to reclaim anyway is very limited. However, step 765 in Figure 7 is performed for an allocation request to the middleware heap, in order to try to reclaim space from the transient heap. The effect of this, if successfully, would generally be to reduce the transient heap below its original size.

As one minor subtlety on the above, in the preferred embodiment, the midpoint position is recalculated when the middleware heap is shrunk (but not when the transient heap size is altered, or when the middleware heap is enlarged), the new position being halfway between the current middleware heap boundary and the current transient heap boundary. This attempts to provide some tuning of the space allocation between the two heaps, although many other algorithms could be considered as the basis for the control procedure.

One complication that arises from effectively having multiple heaps of various sizes is that it becomes more complex to determine whether or not a given object reference is within a heap (as required, for example, for step 612 of Figure 6A), and if so which one (in case, for example, they have different garbage collection policies). One possibility is to compare the reference with the information in the heap control block 530 (see Figure 5). However, with multiple heaps, and also a system heap which is not necessarily contiguous, this becomes a time-consuming operation.

In order to overcome this problem, the preferred embodiment adopts the approach illustrated schematically in Figure 8. As shown, system address space or virtual memory 800 is split into chunks of a standard size, referred to herein as slices 802. As previously mentioned, in the preferred embodiment on a 32 bit system, these slices are each 64KBytes in size. The slices can be numbered linearly as shown with increasing address space. The heaps can then be allocated out of these slices, in such a way that heap space is always allocated or deallocated in terms of an integral number of slices. Figure 8 shows three different heaps (for simplicity termed A , B and C), whereby heap A is non-contiguous and comprises slices 3-4 and 6-7, heap B comprises slice 9, and heap C is contiguous and comprises slices 12-14 inclusive. Note that two or more of these heaps may possibly be being managed as single block of storage (ie in the same manner to the transient and middleware heaps of Figure 5).

Also illustrated in Figure 8 is lookup table 825, which has two columns, the first 830 representing slice number, and the second 831 representing heap number. Thus each row of the table can be used to determine, for the relevant slice, which heap it is in - a value of zero (indicated by a dash) is assumed to indicate that the slice is not currently in a heap. The system updates table 825 whenever slices are allocated to or deallocated from the heap.

Using table 825 it now becomes very quick to determine whether a given memory address is in a heap. Thus an initial determination is made of the relevant slice, by dividing the given memory location (minus the

system base memory location if non-zero) by the slice size, and rounding down to the next integer (ie truncating) to obtain the slice number. This can then be used to directly access the corresponding heap identifier in column 831. In fact, it will be appreciated that column 830 of Table 825 does not need to be stored explicitly, since the memory location of each entry in column 831 is simply a linear function of slice number. More specifically, each entry in column 831 can typically be represented by 1 byte, and so the information for slice N can be found at the base location for table 825, plus N bytes. Overall therefore, this approach provides a rapid mapping from object location to heap identity (if any), irrespective of the number of heaps, or the complexity of their configuration.

One problem however with the technique illustrated in Figure 8 is that on 64 bit machines, the virtual memory or address space is so great that table 825 would become prohibitively large. Thus in a preferred embodiment for such systems, a modified mapping is used, as shown in Figure 9, which has an extra layer in the memory mapping arrangement. In the diagram, memory 900 represents the system address space or virtual memory, which as in Figure 8 is divided into slices 902 (the difference from Figure 8 being that on a 64 bit system, address space is much larger, so there are many more slices). Figure 9 illustrates the location of two heaps, arbitrarily denoted A and B, with A comprising slices 2-4 inclusive, and B comprising slices 1026-1028 inclusive and also slices 9723-9726 inclusive.

Also shown in Figure 9 are two lookup tables, 925, 926, each of which, for the sake of illustration,

contains 2048 entries, and maps to a corresponding range of slices in address space 900. Thus lookup table 925 maps slices 0-2047, whilst lookup table 926 maps slices 8192-10239. These lookup tables are directly analogous to that of Figure 8, in that they logically contain two columns, the first 930 identifying a slice number, and the second 931 the identity of any heap within that slice (or else zero). Tables 925 and 926 can be regarded as forming the lower level of the lookup hierarchy.

Figure 9 also depicts a higher layer in the lookup hierarchy, namely table 940, which again logically contains two columns. The first column 941 logically represents the number of lookup table 925, 926 in the next lower layer of the lookup hierarchy, whilst the second column 942 contains a pointer to the relevant lookup table. Thus the first row of column 942 contains a pointer 951 to table 925, and the fifth row of column 942 contains a pointer 952 to table 926.

It will be noted that to conserve space, lookup tables in the lower level of the hierarchy only exist where at least some of the corresponding slices are assigned to a heap. Thus for the particular arrangement of Figure 9, the lookup tables for slices 2048-4095, 4096-6143, and 6144-8191 have not been created, since none of these slices has been assigned to any heap. In other words, lookup tables 925, 926, etc for various slice ranges will be created and deleted according to whether any slices within that slice range are being utilised for the heap. If this is not the case, and the lookup table is deleted (or not created in the first place), the pointer in column 942 of top level lookup table 940 is set to zero.

The operation of the embodiment shown in Figure 9 is analogous to that of Figure 8, except that there is an extra level of indirection involved in the hierarchy. Thus to determine whether a particular reference or address is within a heap, the correct row is determined based on a knowledge of the size of a slice 902, and also the number of rows in each lower level lookup table 925, 926. It is expected that for most rows, the corresponding entry in column 942 will be null or zero, immediately indicating that that address is not in a heap slice. However, if the lookup selects a row which has a non-zero entry, this is then followed (using pointer 951, 952 or equivalent) to the corresponding lookup table. The desired entry is then found by locating the row using the reference under investigation (allowing for which particular lookup table is involved), and examining the entry for that row in column 931. This will indicate directly whether or not the slice containing the referenced location is in a heap, and if so, which one.

As an example of this, to investigate memory address 637405384 we first integer divide by 65536 (the size of a slice in the preferred embodiment), to give 9727 (truncated), implying we are in slice 9727. Next we perform an integer division of 9727 by 2048 (the number of entries in each lower level look-up table), to give 4 (truncated), implying we are in the 5th row of column 941. It will be appreciated that we could have got here directly by dividing 637405384 by 134217728 (which equals 2048×65536 , or in other words, the total number of addresses per lower level lookup table). In any event, from the 5th row of table 940, it is determined that the corresponding entry in column 941 is non-zero, so that

the specified address may possibly lie in a heap. Accordingly, pointer 952 is followed to table 926. Here we can determine that the row of interest is number 1535 (equal to 9727 modulo 2048), from which we can see that this particular slice is not, after all, part of heap. It follows of course that this is also true for any address within this slice.

Note that as for Figure 8, the slice number columns 930 of lookup tables 925, 926 are not in practice needed, since the desired row in column 931 can be determined directly by using the slice number (modulo 2048) as an offset from the base address of the lookup table. Likewise, column 941 of table 940 is also redundant, since the relevant row can be determined directly from the address. In fact however, the vast majority of rows in table 940 (column 940) are likely to be zero, in which case storing the information in some other data structure such as a linked list would be much more efficient in terms of space (but may reduce lookup speed).

It will be appreciated that any suitable data structure can be used for storing the two levels of lookup information, shown as tables 940, and 925, 926 respectively. It will also be recognised that the sizes discussed with reference to Figures 8 and 9 (a slice size of 65536 bytes; 2048 slices per lower level lookup table) are exemplary only, and can be varied as circumstances dictate to optimize performance.

Returning now to Figure 4, as previously described, at the end of a transaction the transient heap is deleted (equivalent to the refresh heap step 445, performed as part of the reset JVM). This activity is generally similar to garbage collection, although certain

optimizations are possible, and certain additional constraints need to be considered. This process is shown in more detail in the flow chart of Figure 10 (which is split for convenience into two components, 10A and 10B). The first step in Figure 10A (1005) is to wait for all finalization activity to complete. Thus if there has been a GC during a transaction then there may be finalizers to be run and they must be run before the transient heap can be reset, as the finalizers could create (or require) other objects. This checking is performed by confirming that the reference handler and finalizer thread have emptied their respective queues, and that there are no other in-progress objects (ie the processing of all pending finalization objects has been completed). Next all the locks required for garbage collection are obtained, and all other threads are suspended (step 1010). The system is now in a position to commence deletion of the transient heap.

In order to accomplish this, the stacks and registers of all threads are scanned (as for a normal garbage collection), and if a reference is found to the transient heap (step 1015) then the JVM is potentially dirty and so cannot be reset. The reason for this as discussed in relation to standard garbage collection (Figure 6) is that the references on the stacks and registers must be treated as live, even though it is not certain that they are in fact object references. To firm up on this the references are tested to see if it is possible to exclude them from being object references (step 1020), essentially by using the same three tests 612, 615 and 620 of Figure 6. In other words, if the possible reference is not on the heap, or does not fall

on an 8-byte boundary, or does not correspond to an allocated memory location, then it cannot in fact be a reference. Otherwise, the register or stack value may still be a reference, and so processing has to exit with an error that the JVM is dirty and cannot be reset (step 1099). Note that references from the stacks or registers to the middleware or system heap are of course acceptable, because objects on these heaps are not being deleted.

It will be appreciated that based on the above, a spurious data value in a stack or register will sometimes prevent JVM reset. However this happens relatively infrequently in practice, because all but the main application thread and certain system threads should have terminated at this point, so the stacks are relatively empty (nb the policy adopted in the preferred embodiment is that a JVM cannot be reset if more than a single transaction thread was used; multiple middleware threads are tolerated providing they have terminated by the completion of the middleware tidyups). Related to this, as previously mentioned finalizer objects on the transient heap are retained in that heap until a JVM reset. This means that references to such objects are not entered onto the stack for the finalizer thread, which would otherwise typically cause the reset to fail at steps 1015 and 1020 (this would be the case even where the finalize method for the object had been finished, since this would not necessarily lead to complete deletion of the corresponding stack entry; rather the finalizer thread may enter a function to wait for more work, resulting in uninitialized areas on the stack which may point to previously processed finalizer objects).

It is important to note that error 1099 indicating that the JVM is dirty does not imply that previous processing was incorrect, merely that the JVM cannot be reset (although of course this may in turn indicate some unexpected action by the application). In other words, a new JVM will need to be created for the next application. Because of this, if it is detected that the JVM is dirty, such as a negative outcome at step 1020, the method normally proceeds immediately to step 1099. This returns an error code to the reset JVM request from the middleware, with no attempt to continue to perform any further garbage collection. The reason for this is that the middleware may want to do a little more tidying up, but generally it is expected that it will terminate the current JVM fairly quickly. Hence there is unlikely to be a need for any further garbage collection, which rather would represent an unnecessary waste of time. A similar policy is adopted whenever the processing of Figure 10A indicates that the JVM is dirty.

Assuming now a negative result from step 1015 or 1020, the JVM refresh continues with an examination of the primordial statics fields (step 1025) to see what objects they reference. Since these fields will be retained through the JVM reset, it is important that the objects that they reference, either directly or indirectly, are likewise retained. If however the referenced objects are application objects (tested at step 1030) then clearly these cannot be retained, because the application has essentially terminated, and the purpose of resetting the JVM is to allow a new application to commence. Therefore, if the primordial statics do reference an application object, then the JVM

is marked as dirty, and the method proceeds to error 1099.

Assuming that the objects referenced by the primordial static fields are not application objects (typically they will be primordial object instances or arrays), then these are moved ("promoted") from the transient heap to the middleware heap (step 1035). The reason why such objects are placed on the transient heap initially is that at allocation time, it may not be known that the object to be allocated is a primordial static variable, or reachable from one.

(Note that this approach bears some similarities to generational garbage collection, in which new objects are initially allocated to a short-term heap, and then promoted to a longer-term heap if they survive beyond a certain time, but the criterion for promotion is different: essentially it is based on object type or usage, rather than age. Generational garbage collection is discussed further in the book by Jones and Lin referenced above).

One complication (not shown in Figure 10) is that promoting an object from the transient heap to the middleware heap may lead to an allocation failure on the middleware heap if space is exhausted. In such an eventuality, a garbage collection is performed. If this still does not create enough space, then this will lead to error 1099.

After the primordial static objects have been promoted, the next step is to review the card table (536 - see Figure 5). The card table represents a set of bytes, one per fixed unit of heap (for example 512 bytes). Whenever an object reference is written to the

heap, the card table is updated to indicate dirty (nb marking a card as dirty does not imply that the JVM itself is necessarily dirty). The card updated corresponds not to the portion of the heap which contains the updated object reference itself, but rather to the portion of heap which contains the top of the object that includes the the reference (for a small object these may of course be the same). Given that updating object references is a frequent operation, the card table must operate very quickly. This is the reason why each card is a byte despite containing only a single bit of information, because in practice this can be manipulated more quickly. Furthermore, no attempt at this write stage is made to investigate the nature of the reference update, for example whether the reference was set to a null value, or to an object in a particular heap.

Now during JVM reset the card table is scanned, or more particularly those cards which correspond to the region currently assigned to the middleware heap are scanned. Thus cards for the transient heap 520 and for the unassigned region 510 are not scanned, even if they have previously been part of the middleware heap. As part of this review, it is first determined whether any cards are set (ie marked as dirty) (step 1045). This indicates that a reference in the corresponding portion of the middleware heap has been updated since the last JVM reset, and so must be checked to confirm that it does not point to the transient heap. The first part of this check is to find all object references in objects which start in the heap portion corresponding to the marked card. Note that there may be more than one object to review here, or possibly none at all if the object previously

located there has since been garbage collected and the space reused by a larger object whose beginning is situated outside that portion of the heap. For all objects associated with a marked card, all references
5 contained in those objects (even if the references themselves are outside the portion of the heap corresponding to the card) are checked to see if they point to the transient heap (step 1050). If they do not, for example they contain only null pointers, and/or
10 references to the middleware heap, then this is not a problem for JVM reset. On the other hand, if there are any such pointers to the transient heap from the middleware heap, this will be a problem on reset since those references will no longer be valid once the
15 transient heap is cleared. The one exception to this is where the objects containing these problematic references are no longer live (ie could be garbage collected).

Therefore, on a positive outcome to step 1050, the system performs the mark phase of a garbage collection
20 (step 1055), which is a relatively long operation. If the problematic references are in objects which are marked (ie live), as tested at step 1060, then they are indeed problematic, so the JVM must be regarded as dirty; hence the method proceeds to error 1099. On the other hand, if
25 the problematic references are in objects which are not marked, then they can effectively be ignored, since these objects are no longer live.

Note that if the heaps have been compacted during a transaction, then this invalidates the card table. In
30 such cases a full scan of the middleware heap is required to locate an object references to the transient heap,

equivalent to the garbage collection mark phase of step 1055 if any such references are found.

Assuming that the test of step 1060 produces a negative output (ie no live middleware references to the transient heap), the method proceeds to scan JNI global references. These are references which are used by native code routines (ie running directly on OS 30 rather than on JVM 40, see Figure 1) to refer to Java objects. Using the Java Native Interface (JNI) such references can be made global, that is available to all threads, in which case they will exist independently of the thread that created them. All such JNI global reference slots are scanned (step 1065) (see Figure 10B) and if a reference to the transient heap is found (step 1070) the JVM is marked as dirty (ie error 1099), since these references will clearly fail once the transient heap is reset.

Providing this is not the case, the JNI weak references are scanned next (step 1072). These are references which the application specifies using JNI as expendable, in that they can be deleted if no longer used. According, any such weak JNI references to the transient heap that are found can be nulled (step 1074), thereby permitting the JVM reset to proceed.

Next, the static variables of all middleware classes are scanned (step 1076) to see if any directly reference the transient heap (step 1078). Note that these won't previously have been examined, since they are on the system heap rather than the middleware heap. If a direct reference to the transient heap is found, the JVM is dirty, corresponding to error 1099. (Note that unlike for the primordial statics (step 1025) there is no need to iteratively follow references from the middleware

statics, since any indirect references will already have been picked up by preceding analysis). If no transient heap references are found, the processing continues to step 1080 in which objects on the transient heap are reviewed to see if any have finalizer methods., and any that are found are now run (step 1082). One important aspect of the preferred embodiment is that these finalizer methods are run on the main thread, rather than being passed to the system finalizer. An implication of this is that the finalizer methods will be run in the known and controllable context of the main thread. In addition, it is ensured that the finalizer methods complete before progressing to the next stage of the JVM reset. Unfortunately, finalizer methods can create fresh objects, which may newly reference the transient heap. Therefore, after the finalizer methods have completed, processing must return to step 1025 to repeat much of the checking, to ensure that the system is still in a position for JVM reset. In theory, if the finalizer methods have created new objects on the transient heap which themselves have finalizer methods, then this loop may have to be followed more than once.

Note that strictly speaking there is no formal requirement to run the finalizers at this stage, since this is the point at which the JVM would normally terminate at the conclusion of an application, rather than having a garbage collection performed. Nevertheless, the policy in the preferred embodiment is that object finalizers will be run before deletion at JVM reset, although other implementations may have different policies.

It is assumed that eventually all finalizers will be run, resulting in a negative outcome to the test of step 1080. In these circumstances, the method proceeds to step 1085, which represents reset of the JVM by deleting the transient heap. In practice, this involves several operations. Firstly, if the mark phase of the garbage collection was run (step 1055) then the sweep phase, which is relatively quick, is now run on the middleware heap. Next, various operations are performed to formally reset the transient heap, including: the removal of all transient heap monitors and the freeing of storage for transient heap class blocks (ie releasing the storage utilised by the class block, which is not on the heap). The transient heap pointers can now be reset so that the heap is effectively emptied, and restored to its initial size (by setting boundary 522 appropriately).

In the preferred embodiment it is declared that the transient heap will be set to the same initial size for each transaction. One potential problem with honouring this is that the middleware heap may have expanded during the previous application, and then retain this space through a reset of the JVM. Since there is no constraint on the transient heap shrinking below its initial size, to surrender space to the middleware heap if required, this can in turn make it impossible for the transient heap in the next incarnation of the JVM to be set to the same initial size as the current transient heap. If this problem arises, a specific attempt is made to shrink the middleware heap sufficiently to accommodate the correct initial size of the transient heap. However, if this attempt is unsuccessful, the JVM must be marked as dirty, and cannot be reset to its initial state.

Once the transient heap has been recreated (although it could be done before), a garbage collection is performed on the middleware heap if either of the following two cases is true: firstly, if the number of slices left in the unallocated portion of the heap, between the middleware heap and the transient heap, is less than two, or secondly if the amount of free space in the middleware heap plus half the unassigned portion of the heap (see Figure 5) is less than the amount of storage used by the previous transaction times three. Both of these can be regarded as a preemptive garbage collection, performing this operation now if the next transaction is otherwise likely to be constrained for space, in the hope that this will avoid a garbage collection during the transaction itself. Note that in the current implementation this preemptive garbage collection would be performed irrespective of whether a garbage collection mark phase was performed in step 1055. Finally, all the threads can be restarted and the garbage collection locks released, whereupon the reset is completed, and the JVM is available to support the next application.

The skilled person will be aware of many possible variations on the embodiment described above. The invention has been described primarily in relation to Java in a server environment, but it will be understood that it applies to any other language with similar properties (possibly C# from Microsoft Corporation), and is also potentially applicable to the client embodiment, such as when it is necessary to have a quick start-up of applications. In addition, many of the details of the systems and processes utilised are exemplary only, and

can be varied according to particular circumstances. Thus other modifications to the embodiments described herein will be apparent to the skilled person yet remain within the scope of the invention as set out in the attached claims.

5

11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2220
2221